arXiv:1703.06368v1 [cs.PL] 18 Mar 2017

Automating Deductive Verification for Weak-Memory Programs

Alexander J. Summers and Peter Müller

Department of Computer Science, ETH Zurich, Switzerland
{alexander.summers, peter.mueller}@inf.ethz.ch

Abstract. Writing correct programs for weak memory models such as the C11 memory model is challenging because of the weak consistency guarantees these models provide. The first program logics for the verification of such programs have recently been proposed, but their usage has been limited thus far to manual proofs. Automating proofs in these logics via first-order solvers is non-trivial, due to reasoning features such as higher-order assertions, modalities and rich permission resources.
In this paper, we provide the first implementation of a weak memory program logic using existing deductive verification tools. We tackle three recent program logics: Relaxed Separation Logic and two forms of Fenced Separation Logic, and show how these can be encoded using the Viper verification infrastructure. Many ingredients of our encoding are not specific to the particular logics tackled, and we explain the general techniques applied. We evaluate the encoding on examples from existing papers and on code from an open-source library.

1 Introduction

Reasoning about programs running on weak memory is challenging because weak memory models admit executions that are not sequentially consistent, that is, cannot be explained by a sequential interleaving of concurrent threads. Moreover, weak-memory programs employ a range of operations to access memory, which require dedicated reasoning techniques. These operations include fences as well as read and write accesses with varying degrees of synchronisation.

Some of these challenges are addressed by the first program logics for weak-memory programs, in particular, Relaxed Separation Logic (RSL) [28], GPS [26], Fenced Separation Logic (FSL) [13], and FSL++ [12]. These logics apply to interesting classes of C11 programs, but their tool support has been limited to embeddings in Coq. Verification based on these embeddings requires substantial user interaction, which is an obstacle to applying and evaluating these logics.

In this paper, we present a novel approach to automating deductive verification for weak memory programs. We encode large fractions of RSL, FSL, and FSL++ into the intermediate verification language Viper [19], and use the existing Viper verification backends to reason automatically about the encoded programs. This encoding reduces all concurrency and weak-memory features as well as logical

features such as higher-order assertions and custom modalities to a much simpler sequential permission logic.

Our approach is much more lightweight than developing a dedicated verifier for any given weak-memory logic because it utilises Viper's existing verification backends based on symbolic execution or verification condition generation. On the other hand, our approach achieves a higher degree of automation than Coq, without requiring tactics that are specific to a particular logic. Saving the effort to develop dedicated backends or tactics is especially useful for prototyping and evaluating logics.

The contributions of this paper are: (1) The first automated verification approach for weak-memory logics. We demonstrate the effectiveness of this approach on examples from the literature, which are available online [3]. (2) An encoding of large fractions of RSL, FSL, and FSL++ into Viper. Various aspects of this encoding (such as the treatment of higher-order features and modalities, as well as the overall proof search strategy) are generic and can be reused to encode other advanced separation logics. In the following, we will refer to RSL, FSL, and FSL++ collectively as *the RSL logics*.

Related Work. The existing weak-memory logics RSL [28], GPS [26], FSL [13], and FSL++ [12] have been formalized in Coq and used to verify small examples. The proofs were constructed mostly manually, whereas our approach automates most of the proof steps. The degree of automation in Coq could be increased through logic-specific tactics [24,8], whereas our approach benefits from Viper's automation for the intermediate language, which is independent of the encoded logic. On the other hand, the Coq formalizations are foundational whereas the soundness of our verification results depends on the correctness of the Viper tool. While foundational proofs are important for critical applications, a more lightweight approach is useful especially to gain experience with new logics.

Jacobs [15] proposed a program logic for the TSO memory model that has been encoded in VeriFast. However, verification in this encoding requires a substantial amount of annotations, whereas our approach provides a higher degree of automation and handles the more complex C11 memory model. Caper [11] is a verifier targeting programs with fine-grained concurrency; the implemented proof technique does not, however, tackle weak memory reasoning.

Automating logics by an encoding into intermediate verification languages has proven useful for sequential programs, as witnessed by the many existing verifiers (e.g. [17,16,10,4,9]) which target Boogie [4] or Why3 [5]. Our work is the first that applies this approach to logics for weak-memory concurrency.

Outline. The next four sections present the encoding for the core features of the C11 memory model: we discuss non-atomic locations in Sec. 2, release-acquire accesses in Sec. 3, fences in Sec. 4, and compare and swap in Sec. 5. We introduce the necessary background on Viper along the way. We discuss additional features in Sec. 6 and evaluate our approach in Sec. 7. Sec. 8 concludes. The appendix

$$s ::= l := \mathtt{alloc}_{\mathsf{na}}() \mid l := \mathtt{alloc}_{\rho}(\mathcal{Q}) \mid [l]_{\sigma} := e \mid x := [l]_{\sigma}$$
$$\mid \mathtt{fence}_{\mathsf{acq}} \mid \mathtt{fence}_{\mathsf{rel}}(A) \mid x := \mathtt{CAS}_{\tau}(l, e_1, e_2)$$
$$\textit{where } \rho \in \{\mathtt{acq}, \mathtt{RMW}\}, \quad \sigma ::= \mathtt{na} \mid \tau, \quad \tau \in \{\mathtt{acq}, \mathtt{rel}, \mathtt{rel_acq}, \mathtt{rlx}\}$$

Fig. 1. Syntax for memory accesses. na indicates a non-atomic operation; τ indicates an atomic access mode (as defined in C11), discussed in later sections. ρ, and assertions A and invariants $\mathcal{Q}$ are program annotations, needed as input for our encoding. Expressions e include boolean and arithmetic operations, but no heap accesses. We assume that source programs are type-checked.

$$A ::= e \mid l \overset{k}{\mapsto} e \mid A_1 * A_2 \mid e \Rightarrow A \mid (e \ ?\ A_1 : A_2)$$
$$\mid \mathsf{Uninit}(l) \mid \mathsf{Acq}(l, \mathcal{Q}) \mid \mathsf{Rel}(l, \mathcal{Q}) \mid \mathsf{Init}(l) \mid \triangle A \mid \bigtriangledown A \mid \mathsf{RMWAcq}(l, \mathcal{Q})$$

Fig. 2. Assertion syntax of the RSL logics. The top row of constructs are standard for separation logics; those in the second row are specific to the RSL logics, and explained throughout the paper. Invariants $\mathcal{Q}$ are *functions* from values to assertions (cf. Sec. 3). We replace RSL's emp assertion by true: see Sec. 6.

contains further details on the encoding and our examples. The Viper code for all examples is available online [3].

2 Non-atomic Locations

We present our encoding for a small imperative programming language that reflects the core operations relevant for C-like programs using C11 atomics, and is similar to the languages supported by the RSL logics. C11 supports *non-atomic* memory accesses and different forms of *atomic* accesses. The access operations are summarised in Fig. 1. We adopt the common simplifying assumption [28,26] that memory locations are partitioned into those accessed only via non-atomic accesses (*non-atomic locations*), and those accessed only via C11 atomics (*atomic locations*). Read and write statements are parameterised by a mode σ, which is either na (non-atomic) or one of the atomic access modes[1] τ. We focus on non-atomic accesses in this section and discuss atomics in subsequent sections.

RSL proof rules. Non-atomic memory accesses come with no synchronisation guarantees; programmers are obliged to ensure that all accesses to non-atomic locations are data-race free. The RSL logics enforce this requirement using standard separation logic [20,23]. We show the syntax of assertions in Fig. 2, which will be explained throughout the paper. A *points-to assertion* $l \overset{k}{\mapsto} e$ denotes a transferrable *resource*, providing permission to access the location l, and expressing that the location has been initialised and its current value is the

[1] C11 provides further "consume" and "sequentially consistent" access modes. The former is not supported by the RSL logics, while the latter has the same proof rules as for the rel_acq mode; both are omitted here, for brevity.

$$\frac{}{\vdash \{\mathsf{true}\}\ l := \mathtt{alloc}_{\mathsf{na}}()\ \{\mathsf{Uninit}(l)\}} \qquad \frac{}{\vdash \{l \overset{1}{\mapsto} _ \lor \mathsf{Uninit}(l)\}\ [l]_{\mathsf{na}} := e\ \{l \overset{1}{\mapsto} e\}}$$

$$\frac{}{\vdash \{l \overset{k}{\mapsto} e\}\ x := [l]_{\mathsf{na}}\ \{x = e * l \overset{k}{\mapsto} e\}} \qquad (l \overset{k}{\mapsto} e\ *\ l \overset{k'}{\mapsto} e') \Leftrightarrow (e = e' * l \overset{k+k'}{\mapsto} e)$$

Fig. 3. Adapted RSL rules for non-atomics. Read access requires a non-zero permission. Write access requires either write permission or that the location is uninitialised. The underscore _ stands for an arbitrary value.

value of expression e. Here, k is a fraction $0 < k \leq 1$; $k = 1$ denotes the *full* (or exclusive) permission to read and write location l, whereas $0 < k < 1$ provides (non-exclusive) read access [7]. Points-to resources can be split and recombined, but never duplicated or forged; when transferring such a resource to another thread it is removed from the current one, avoiding data races by construction. The RSL assertion $\mathsf{Uninit}(l)$ expresses exclusive access to a location l that has been allocated, but not yet initialised; l may be written to but not read from (in C11 this would be undefined behaviour). The main proof rules for non-atomic locations, adapted from RSL [28], are shown in Fig. 3.

Encoding. The Viper intermediate verification language [19], which we use to encode the RSL logics, supports an assertion language based on Implicit Dynamic Frames [25], a program logic related to separation logic [21], but which separates permissions from value information. Viper is object-based; the only memory locations are field locations $e.f$ (in which e is a reference, and f a field name). Permissions to access these heap locations are described by *accessibility predicates* of the form `acc(`$e.f$`, `k`)`, where k is a fraction as for points-to predicates above (k defaults to 1). Assertions that do not contain accessibility predicates are called *pure*. Unlike in separation logics, heap locations may be read in pure assertions.

We model C-like memory locations l using a field `val` of a Viper reference l. Consequently, a separation logic assertion $l \overset{k}{\mapsto} e$ is represented in Viper as `acc(`l`.val, `k`) && `l`.val == e`. In this paper, we assume that memory locations have type `int`, but a generalisation is trivial. Viper's conjunction `&&` treats permissions multiplicatively, requiring the sum of the permissions in each conjunct, and acts as logical conjunction for pure assertions (just as $*$ in separation logic).

Viper provides two key statements for encoding proof rules. An *inhale* statement **`inhale`** A has the effect of *adding* the permissions denoted by the Viper assertion A to the current state, and *assuming* pure assertions in A. This can be used to model gaining new resources, e.g., acquiring a lock in the source program. Dually, **`exhale`** A checks that the current state satisfies A (otherwise a verification error occurs), and *removes* the permissions that A denotes; the values of any locations to which no permission remains are *havoced* (assigned arbitrary values). For example, when forking a new thread, its precondition is exhaled to reflect transferring ownership of the necessary resources from the forking thread. Inhale

```
field val:  Int
field init: Bool

⟦Uninit(l)⟧ ⇝ acc(l.val) && acc(l.init) && !l.init
⟦l ↦^k e⟧ ⇝ acc(l.val, k) && acc(l.init, k) && l.val == ⟦e⟧ && l.init

⟦l := alloc_na()⟧ ⇝ l := new(); inhale ⟦Uninit(l)⟧
⟦x := [l]_na⟧ ⇝ assert l.init; x := l.val
⟦[l]_na := e⟧ ⇝ l.val := ⟦e⟧; l.init := true
```

Fig. 4. Viper encoding of the RSL assertions and rules for non-atomic memory accesses.

and exhale statements can be seen as the permission-aware analogues of the assume and assert statements of first-order verification languages [17].

The encoding of the rules for non-atomics from Fig. 3 is presented in Fig. 4. $\llbracket A \rrbracket \rightsquigarrow \ldots$ denotes the encoding of an RSL assertion A as a Viper assertion, and analogously $[\![s]\!] \rightsquigarrow \ldots$ for source-level statements s.

The first two lines show background declarations employed in our encoding. The assertion encodings follow the explanations above. Allocation is modelled by obtaining a fresh reference (via `new()`) and inhaling permissions to its `val` and `init` fields; Assuming `!l.init` reflects that the location is not yet initialised. Viper implicitly checks the necessary permissions for field accesses (verification fails otherwise). Hence, the translation of a non-atomic reads only needs to check that the read location is initialised before obtaining its value. Analogously, because of Viper's implicit permission check, the translation of a non-atomic write only stores the value and records that the location is now initialised.

Note that Viper's implicit permission checks are both necessary and sufficient to encode the RSL rules in Fig. 3. In particular, the assertions $l \overset{1}{\mapsto} _$ and $\mathsf{Uninit}(l)$ both provide the permissions to write to location l. By including `acc(`l`.val)` in the encoding of both assertions, we avoid the disjunction of the RSL rule[2].

Our approach requires programmers to annotate their code with access modes for locations (as part of the `alloc` statement), and specifications such as pre and postconditions for methods and threads. Given these inputs, Viper constructs the proof automatically. In particular, it automatically proves entailments, and splits and combines fractional permissions (hence, the equivalence in Fig. 3 need not be encoded). Automation can be increased further by inferring some of the required assertions, but this is orthogonal to the encoding presented in this paper.

3 Release-Acquire Atomics

The simplest form of C11 atomic memory accesses are *release write* and *acquire read* operations. They can be used to synchronise the transfer of ownership of

[2] In fact, Viper does not support disjunction over impure assertions. This does not restrict expressiveness in practice, since disjunctions used in RSL specifications typically include a pure condition which disambiguates the cases such that they can be encoded using implications and conditional expressions.

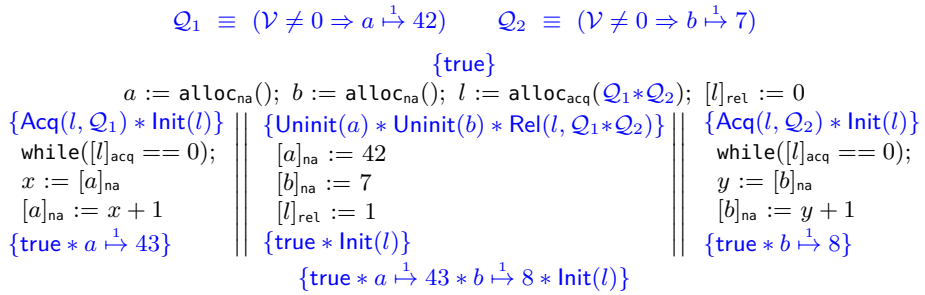

Fig. 5. An example illustrating "message passing" of non-atomic ownership, using release acquire atomics (inspired by an example from [13]); annotations shown in blue.

(and information about) other, non-atomic locations, using a *message passing idiom*, illustrated by the example in Fig. 5. This program allocates two non-atomic locations a and b, and an atomic location l (initialised to 0), which is used to synchronise the three threads that are spawned afterwards. The middle thread makes changes to the non-atomics a and b, and then signals completion via a release write of 1 to l; conceptually, it gives up ownership of the non-atomic locations via this signal. The other threads loop attempting to acquire read a non-zero value from l. Once they do, they each gain ownership of one non-atomic location via the acquire read of 1 and access that location. The release write and acquire reads of value 1 enforce *ordering constraints* on the non-atomic accesses, preventing the left and right threads from racing with the middle one.

RSL proof rules. The RSL logics capture message-passing idioms by associating a *location invariant* $\mathcal{Q}$ with each atomic location. Such an invariant is a function from values to assertions; we represent such functions as assertions with a distinguished variable symbol $\mathcal{V}$ as parameter. Location invariants prescribe the intended ownership that a thread obtains when performing an acquire read of value $\mathcal{V}$ from the location, and that must correspondingly be given up by a thread performing a release write. The main proof rules [28] are shown in Fig. 6.

When allocating an atomic location for release/acquire accesses (first proof rule), a location invariant $\mathcal{Q}$ must be chosen (as an annotation on the allocation). The assertions $\mathsf{Rel}(l, \mathcal{Q})$ and $\mathsf{Acq}(l, \mathcal{Q})$ record the invariant to be used with subsequent release writes and acquire reads. To perform a release write of value e (second rule), a thread must hold the $\mathsf{Rel}(l, \mathcal{Q})$ assertion and *give up* the assertion $\mathcal{Q}[e/\mathcal{V}]$. For example, the line $[l]_{\mathsf{rel}} := 1$ in Fig. 5 causes the middle thread to give up ownership of both non-atomic locations a and b. The assertion $\mathsf{Init}(l)$ represents that atomic location l is initialised; both $\mathsf{Init}(l)$ and $\mathsf{Rel}(l, \mathcal{Q})$ are *duplicable* assertions; once obtained, they can be passed to multiple threads.

Multiple acquire reads might read the value written by a single release write operation; RSL prevents ownership of the transferred resources from being obtained (unsoundly) by multiple readers in two ways. First, $\mathsf{Acq}(l, \mathcal{Q})$ assertions

cannot be duplicated, only split by partitioning the invariant $\mathcal{Q}$ into disjoint parts. For example, in Fig. 5, $\mathsf{Acq}(l, \mathcal{Q}_1)$ is given to the left thread, and $\mathsf{Acq}(l, \mathcal{Q}_2)$ to the right. Second, the rule for acquire reads adjusts the invariant in the Acq assertion such that subsequent reads of the same value will not obtain any ownership.

Encoding. A key challenge for encoding the above proof rules is that Rel and Acq are parameterised by the invariant $\mathcal{Q}$; higher-order assertions are not directly supported in Viper. However, for a given program, only finitely many such parameterisations will be required, which allows us to apply defunctionalisation [22], as follows. Given an annotated program, we assign a unique *index* to each syntactically-occurring invariant $\mathcal{Q}$ (in particular, in allocation statements, and as parameters to Rel and Acq assertions in specifications). Furthermore, we assign unique indices to all *immediate conjuncts* of these invariants. For the program in Fig. 5, we assign indices 1, 2 and 3 to invariants $\mathcal{Q}_1$, $\mathcal{Q}_2$ and $\mathcal{Q}_1 * \mathcal{Q}_2$, respectively. We write *indices* for the set of indices used. For each i in *indices*, we write $inv(i)$ for the invariant which i indexes. For an invariant $\mathcal{Q}$ from the source program, we write $\langle\mathcal{Q}\rangle$ for its index, and $\langle\!\langle\mathcal{Q}\rangle\!\rangle$ for the set of indices assigned to its immediate conjuncts. In our example, $\langle\mathcal{Q}_1 * \mathcal{Q}_2\rangle$ is 3, while $\langle\!\langle\mathcal{Q}_1 * \mathcal{Q}_2\rangle\!\rangle$ is $\{1, 2\}$.

Our encoding of the RSL rules from Fig. 6 is summarised in Fig. 7. To encode duplicable assertions such as $\mathsf{Init}(l)$, we make use of Viper's *wildcard permissions* [19], which represent unknown positive permission amounts. When exhaled, these amounts are *angelically chosen* [14]; provided *some* positive permission is held, the amount exhaled will always be *strictly smaller*. So after inhaling an $\mathsf{Init}(l)$ assertion (that is, a **`wildcard`** permission), it is possible to exhale two **`wildcard`** permissions, corresponding to two $\mathsf{Init}(l)$ assertions. Note that for atomic locations, we only use the `init` field's permissions, not its value.

We represent a $\mathsf{Rel}(l, _)$ assertion for *some* invariant via a **`wildcard`** permission to a `rel` field; this is represented via the `SomeRel(l)` macro[3], and is used as the precondition for a release write (we must hold *some* Rel assertion, according to Fig. 6). The specific invariant associated with the location l is represented by

[3] Viper macros can be defined for assertions or statements, and are syntactically expanded (and their arguments substituted) on use.

$$\frac{}{\vdash \{\mathsf{true}\}\ l := \mathtt{alloc}_{\mathtt{acq}}(\mathcal{Q})\ \{\mathsf{Rel}(l, \mathcal{Q}) * \mathsf{Acq}(l, \mathcal{Q})\}}$$

$$\frac{}{\vdash \{\mathcal{Q}(e) * \mathsf{Rel}(l, \mathcal{Q})\}\ [l]_{\mathtt{rel}} := e\ \{\mathsf{Init}(l) * \mathsf{Rel}(l, \mathcal{Q})\}}$$

$$\frac{}{\vdash \{\mathsf{Init}(l) * \mathsf{Acq}(l, \mathcal{Q})\}\ x := [l]_{\mathtt{acq}}\ \{\mathcal{Q}[x/\mathcal{V}] * \mathsf{Acq}(l, \mathcal{V} \neq x \Rightarrow \mathcal{Q})\}}$$

$$\mathsf{Init}(l) \ \Leftrightarrow\ \mathsf{Init}(l) * \mathsf{Init}(l) \qquad \mathsf{Rel}(l, \mathcal{Q}) \ \Leftrightarrow\ \mathsf{Rel}(l, \mathcal{Q}) * \mathsf{Rel}(l, \mathcal{Q})$$

$$\mathsf{Acq}(l, \mathcal{Q}_1 * \mathcal{Q}_2) \ \Leftrightarrow\ \mathsf{Acq}(l, \mathcal{Q}_1) * \mathsf{Acq}(l, \mathcal{Q}_2) \qquad \mathcal{Q}_1 \models \mathcal{Q}_2 \ \Rightarrow\ \mathsf{Acq}(l, \mathcal{Q}_1) \models \mathsf{Acq}(l, \mathcal{Q}_2)$$

Fig. 6. Adapted RSL rules for release-acquire atomics.

```
field rel: Int
field acq: Bool
predicate AcqConjunct(l: Ref, idx: Int)

function valsRead(l: Ref, i: Int): Set[Int]
  requires AcqConjunct(l, i)

define SomeRel(l)  acc(l.rel, wildcard)
define SomeAcq(l)  acc(l.acq, wildcard) && l.acq == true

⟦Init(l)⟧ ⇝ acc(l.init, wildcard)
⟦Rel(l, Q)⟧ ⇝ SomeRel(l) && l.rel == ⟨Q⟩
⟦Acq(l, Q)⟧ ⇝ SomeAcq(l) && (foreach i in ⟪Q⟫:
  AcqConjunct(l, i) && valsRead(l, i) == Set[Int]() end)

⟦l := alloc_acq(Q)⟧ ⇝ l := new(); inhale ⟦Rel(l, Q)⟧ && ⟦Acq(l, Q)⟧

⟦[l]_rel := e⟧ ⇝ assert SomeRel(l);
  foreach i in indices do
    if (i == l.rel) { exhale inv(i)[e/V] }
  end
  inhale Init(l)

⟦x := [l]_acq⟧ ⇝ assert Init(l) && SomeAcq(l); x := havoc(); // unknown Int
  foreach i in indices do
    if (perm(AcqConjunct(l, i)) == 1 && !(x in valsRead(l, i))) {
      inhale inv(i)[x/V]
      tmpSet := valsRead(l, i)
      exhale AcqConjunct(l, i)
      inhale AcqConjunct(l, i) && valsRead(l,i) == tmpSet union Set(x)
    }
  end
```

Fig. 7. Viper encoding of the RSL rules for release-acquire atomics. The operations in italics (in particular, the *foreach*...*end* repetitions) are expanded statically in our encoding (into conjunctions or statement sequences). The (here, always **true**) value of the acq field will be explained in Sec. 5.

storing its index as the *value* of the rel field; when encoding a release write, we branch on this value to exhale the appropriate assertion[4].

Analogously to Rel, we represent an Acq assertion for *some* invariant using a **wildcard** permission (the SomeAcq macro), which is the precondition for executing an acquire read. However, to support splitting, we represent the invariant in a more fine-grained way, by recording individual conjuncts separately. Each conjunct i of the invariant is modelled as an abstract *predicate* instance AcqConjunct(l, i), which can be inhaled and exhaled individually, just like individual field permissions. This encoding handles the common case that invariants are split along top-level conjuncts, as in Fig. 5. More complex splits can be supported through additional annotations, see App. C.

[4] The RSL logics provide rules for rewriting and splitting the invariants of Rel assertions; these can be used to make the invariant stronger than necessary, which is not useful in practice [27], and thus not considered here.

A release write is encoded by checking that some Rel assertion is held, and then exhaling the associated invariant for the value written. Moreover, it records that the location is initialised. The RSL rule for acquire reads adjusts the Acq invariant by obliterating the assertion for the value read. Instead of directly representing the adjusted invariant (which would complicate our numbering scheme), we track the set of values read as state in our encoding. We complement each `AcqConjunct` predicate instance with an (uninterpreted) Viper function `valsRead(`l`, `i`)` depending on it, returning a set of indices[5].

A require read checks that the location is initialised and that we have *some* Acq assertion for the location. It assigns an unknown value to the lhs variable x, which is subsequently constrained by the invariant associated with the Acq assertion as follows: We check for each index whether we both currently hold an `AcqConjunct` predicate for that index[6], and if so, have not previously read the value x from that conjunct of our invariant. If these checks succeed, we inhale the indexed invariant for x, and then include x in the values read.

The encoding presented so far allows us to automatically verify annotated C11 programs using release writes and acquire reads (e.g., the program of Fig. 5) without any custom proof strategies [3]. In particular, we can support the higher-order Acq and Rel assertions through defunctionalisation and enable the splitting of invariants through a suitable representation.

4 Relaxed Memory Accesses and Fences

In contrast to release-acquire accesses, C11's *relaxed* atomic accesses provide no synchronisation. Consequently, RSL's proof rules for relaxed atomics provide weak guarantees, in particular, do not support ownership transfer.

Memory fence instructions can eliminate this problem. Intuitively, a *release fence* together with a subsequent relaxed write allows a thread to transfer away ownership of resources, similar to a release write (provided that the transferred resources are not accessed between the fence and the relaxed write). Dually, an *acquire fence* together with a prior relaxed read allows a thread to obtain ownership of resources, similar to an acquire read (again, provided that the transferred resources are not accessed between the relaxed read and the fence); this reasoning is justified by the ordering guarantees of the C11 model [13].

FSL proof rules. FSL and FSL++ provide proof rules for fences (see Fig. 8). Theyuse *modalities* $\triangle$ ("up") and $\triangledown$ ("down") to represent resources that are transferred through relaxed accesses and fences. An assertion $\triangle A$ represents a resource A which has been prepared, via a release fence, to be transferred by a relaxed write operation; dually, $\triangledown A$ represents resources A obtained via a relaxed read, which may not be made use of until an acquire fence is encountered. The

[5] Viper's heap-dependent functions are mathematical functions of their parameters and the resources stated in their preconditions (here, `AcqConjunct(`l`,`i`)`).

[6] A **`perm`** expression yields the permission fraction held for a field or predicate instance.

$$\frac{}{\{A\}\ \texttt{fence}_{\texttt{rel}}\ \{\triangle A\}} \qquad \frac{}{\{\triangledown A\}\ \texttt{fence}_{\texttt{acq}}\ \{A\}}$$

$$\frac{}{\{\triangle\mathcal{Q}(e) * \mathsf{Rel}(l,\mathcal{Q})\}\ [l]_{\texttt{rlx}} := e\ \{\mathsf{Init}(l) * \mathsf{Rel}(l,\mathcal{Q})\}}$$

$$\frac{}{\{\mathsf{Init}(l) * \mathsf{Acq}(l,\mathcal{Q})\}\ x := [l]_{\texttt{rlx}}\ \{\triangledown\mathcal{Q}[x/\mathcal{V}] * \mathsf{Acq}(l, \mathcal{V} \neq x \Rightarrow \mathcal{Q})\}}$$

$$(A_1 \Rightarrow A_2) \Leftrightarrow (\triangle A_1 \Rightarrow \triangle A_2) \Leftrightarrow (\triangledown A_1 \Rightarrow \triangledown A_2)$$

$$\triangledown(A_1 * A_2) \equiv (\triangledown A_1) * (\triangledown A_2) \textit{ and analogously for } \triangle \textit{ and other binary connectives}$$

Fig. 8. Adapted FSL rules for relaxed atomics and fences.

proof rule for relaxed write is identical to that for a release write (cf. Fig. 6), except that the assertion to be transferred away must be under the $\triangle$ modality; this can be achieved by the rule for release fences. The rule for a relaxed read is the same as that for acquire reads, except that the gained assertion is under the $\triangledown$ modality. The modality can be removed by a subsequent acquire fence. Finally, assertions may be rewritten under modalities, and both modalities distribute over all other logical connectives.

Encoding. The main challenge in encoding the FSL rules for fences is how to represent the two new modalities. Since these modalities guard assertions which cannot be currently used or combined with modality-free assertions, we model them using two *additional heaps* to represent the assertions under each modality. The program heap (along with associated permissions) is a built-in notion in Viper, and so we cannot directly employ three heaps. Therefore, we construct the additional "up" and "down" heaps, by axiomatising bijective mappings `up` and `down` between a real program reference and its counterparts in these heaps. That is, technically our encoding represents each source location through three references in Viper's heap (rather than one reference in three heaps). Assertions $\triangle A$ are then represented by replacing *all references* r in the encoded assertion A with their counterpart $\texttt{up}(r)$. We write $\lceil A \rceil^{up}$ for the transformation which performs this replacement. For example, $\lceil$`acc(x.val) && x.val == 4`$\rceil^{up} \rightsquigarrow$ `acc(up(x).val) && up(x).val == 4`. We write $\lceil A \rceil^{down}$ for the analogous transformation for the `down` function.

The extension of our encoding is shown in Fig. 9. We employ a Viper *domain* to introduce and axiomatise the mathematical functions used to represent our `up` and `down` mappings. By axiomatising inverses for these mappings, we guarantee bijectivity. Bijectivity allows Viper to conclude that (dis)equalities and other information is preserved under these mappings. Consequently, we do not have to explicitly encode the last two rules of Fig. 8; they are reduced to standard assertion manipulations in our encoding.

An additional `heap` function labels references with an integer identifying the heap to which they belong (we use `0` for real references, and `-1` and `1` for their

```
domain threeHeaps {
  function up(x: Ref) : Ref;    function upInv(x: Ref) : Ref;
  function down(x: Ref) : Ref;  function downInv(x: Ref) : Ref;
  function heap(x: Ref) : Int; // identifies which heap a Ref is from

  axiom { forall r:Ref :: upInv(up(r)) == r &&
    (heap(r) == 0 ==> heap(up(r)) == 1 }
  axiom { forall r:Ref :: up(upInv(r)) == r &&
    (heap(r) == 1 ==> heap(upInv(r)) == 0 }
  axiom { forall r:Ref :: downInv(down(r)) == r &&
    (heap(r) == 0 ==> heap(down(r)) == -1 }
  axiom { forall r:Ref :: down(downInv(r)) == r &&
    (heap(r) == -1 ==> heap(downInv(r)) == 0 }
}
```

$[\![\triangle A]\!] \rightsquigarrow \lceil [\![A]\!] \rceil^{up}$ $\quad$ $[\![\triangledown A]\!] \rightsquigarrow \lceil [\![A]\!] \rceil^{down}$

$[\![[l]_{\mathtt{rlx}} := e]\!] \rightsquigarrow$ *...encoded as for release writes (Fig. 7) except using* $\lceil inv(i) \rceil^{up}$ *in place of* $inv(i)$

$[\![x := [l]_{\mathtt{rlx}}]\!] \rightsquigarrow$ *...encoded as for acquire reads (Fig. 7) except using* $\lceil inv(i) \rceil^{down}$ *in place of* $inv(i)$

$[\![\mathtt{fence}_{\mathtt{rel}}(A)]\!] \rightsquigarrow$ **exhale** $[\![A]\!]$; **inhale** $\lceil [\![A]\!] \rceil^{up}$

$[\![\mathtt{fence}_{\mathtt{acq}}]\!] \rightsquigarrow$

```
var rs : Set[Ref]; rs := havoc() // unknown set of Refs
assume forall r: Ref :: r in rs <==> perm(down(r).val) > none
inhale forall r: Ref :: r in rs ==> acc(r.val, perm(down(r).val))
assume forall r: Ref :: r in rs ==> r.val == down(r).val
exhale forall r: Ref :: r in rs ==> acc(down(r).val, perm(down(r).val))
// analogously for each other field, predicate (in place of val)
```

Fig. 9. Viper encoding of the FSL rules for relaxed atomics and memory fences. We omit triggers for the quantifiers for simplicity, but see [3].

"down" and "up" counterparts, respectively); this labelling provides the verifiers with the (important) information that these notional heaps are disjoint.

Our handling of relaxed read and write is almost identical to that of acquire reads and release writes in Fig. 7; this similarity comes from the original proof rules, which only require that the modalities be inserted when dealing with the invariant. Our encoding for release fences requires an annotation in the source program to indicate which assertion to prepare for release by placing it under the $\triangle$ modality. Inferring these assertions automatically is future work.

For the dual case of acquire fences, we can provide an encoding which does *not* require any user annotations. *Any* assertion that is under the $\triangledown$ modality can (and should) be converted to its corresponding version without the modality because $\triangledown A$ is strictly less-useful than A itself. We encode this conversion as follows. We find *all* permissions currently held in the down heap, and transfer the permissions to (and values of) these locations over to the real heap. We encode this conversion for each field and predicate separately; the steps for the `val` field are shown at the end of Fig. 9. We first define a set `rs` to be precisely the set of all references `r` to which *some* permission to the field location `down(r).val` is currently held, i.e., **`perm`**`(down(r).val) > none`. Then, for each such location we **`inhale`** exactly the same amount of permission to the corresponding `r.val` location.

$$\frac{}{\{\mathsf{true}\}\ l := \mathtt{alloc}_{\mathsf{RMW}}(\mathcal{Q})\ \{\mathsf{Rel}(l,\mathcal{Q}) * \mathsf{RMWAcq}(l,\mathcal{Q})\}}$$

$$\frac{\begin{array}{ll} x \notin FV(P) & P' \equiv \begin{cases} P & \text{if } \tau \in \{\mathtt{rel}, \mathtt{rel_acq}\} \\ \triangle P & \text{otherwise} \end{cases} \\ x \notin FV(e) & \\ \mathcal{Q}[e/\mathcal{V}] \models A * T & A' \equiv \begin{cases} A & \text{if } \tau \in \{\mathtt{acq}, \mathtt{rel_acq}\} \\ \triangledown A & \text{otherwise} \end{cases} \\ P * T \models \mathcal{Q}[e'/\mathcal{V}] & \end{array}}{\left\{\begin{array}{l} \mathsf{Init}(l) * \mathsf{Rel}(l,\mathcal{Q}) * \\ \mathsf{RMWAcq}(l,\mathcal{Q}) * P' \end{array}\right\}\ x := \mathtt{CAS}_\tau(l,e,e')\ \left\{\begin{array}{l} (x = e\ ?\ A' : P') * \mathsf{Init}(l) * \\ \mathsf{Rel}(l,\mathcal{Q}) * \mathsf{RMWAcq}(l,\mathcal{Q}) \end{array}\right\}}$$

$$\mathsf{RMWAcq}(l,\mathcal{Q}) \quad \Leftrightarrow \quad \mathsf{RMWAcq}(l,\mathcal{Q}) * \mathsf{RMWAcq}(l,\mathcal{Q})$$

Fig. 10. Adapted FSL++ rules for compare and swap operations. FV yields the free variables of an assertion.

Viper allows such pointwise updates to the permissions held, via its support for permissions under quantifiers [19,18]. Having added permission to these locations, we can equate the heap values, and then remove the permission to the down locations, completing the conversion of these locations.

With our encoding based on multiple heaps, reasoning about assertions under modalities automatically inherits all of the natively-supported automation that Viper provides for permission and heap reasoning. We will reuse this idea for a different purpose in the following section.

5 Compare and Swap

C11 includes atomic *read-modify-write* operations, commonly used to implement high-level synchronisation primitives such as locks. In particular, the FSL++ logic [12] provides proof rules for *compare and swap* (CAS) operations, as supported natively by many architectures and libraries. An atomic compare and swap $\mathtt{CAS}_\tau(l,e,e')$ reads and returns the value of location l; if the value read is equal to e, it also writes the value e' (otherwise we say that the CAS *fails*).

FSL++ proof rules. The latest proof rules for C11 CAS operations come from FSL++. They extend the previously-presented logic with a new assertion $\mathsf{RMWAcq}(l,\mathcal{Q})$, which is similar to $\mathsf{Acq}(l,\mathcal{Q})$, but is used for transferring ownership via CAS operations instead of acquire reads. The essential idea is for a successful CAS operation to *both* obtain ownership of an assertion via its read operation, and give up ownership of an assertion via its write operation.

FSL++ does not support general combinations of atomic reads and CAS operations on the same location; the way of reading must be chosen at allocation via the annotation ρ on the allocation statement (see Fig. 1). In contrast to the Acq assertions used for atomic reads, RMWAcq assertions can be freely duplicated and their invariants need not be rewritten for a successful CAS: when using only CAS operations, each value read from a location corresponds to a different write.

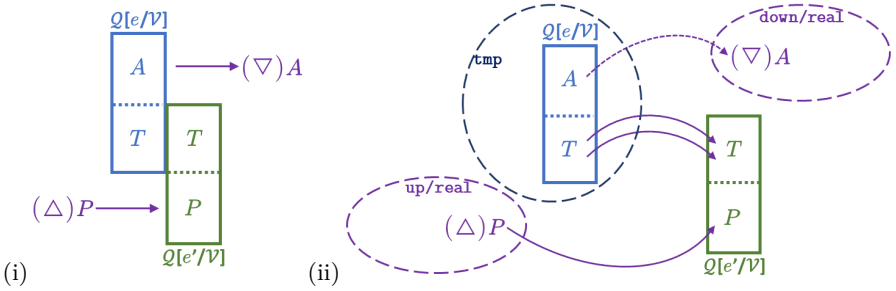

Fig. 11. (i) An informal illustration of the proof rule for CAS operations. (ii) An illustration of our Viper encoding of the proof rule; the dashed regions denote the relevant heaps employed in the encoding.

Our presentation of the relevant proof rules is shown in Fig. 10. Allocating a location with annotation `RMW` provides a Rel and a RMWAcq assertion, such that the location can be used for release writes and CAS operations.

For the CAS operation, we present a single, general proof rule instead of four rules for the different combinations of access modes in FSL++[7]. The rule requires that l is initialised (since its value is read), Rel and RMWAcq assertions (the fact that Rel is required is a technicality, which we will not explain here), and an assertion P' that provides the resources needed for a successful CAS. If the CAS fails (that is, $x \neq e$), its precondition is preserved.

If the CAS succeeds, it has read value e and written value e'. Assuming for now that the access mode τ permits ownership transfer, the thread has acquired $\mathcal{Q}[e/\mathcal{V}]$ and released $\mathcal{Q}[e'/\mathcal{V}]$. As illustrated in Fig. 11(i), these assertions may overlap. Let T denote the assertion characterizing the overlap; then assertions A denotes $\mathcal{Q}[e/\mathcal{V}]$ without the overlap, and P denotes $\mathcal{Q}[e'/\mathcal{V}]$ without the overlap. The net effect of a successful CAS is then to acquire A and to release P, whereas T remains with the location invariant across the CAS. The proof rule is phrased in terms of T, A, and P. Automating the choice of these parameters is one of the main challenges of encoding this rule. Finally, if the access mode τ does not permit ownership transfer (that is, fences are needed to perform the transfer), A and P are put under the appropriate modalities.

Encoding. Our encoding of CAS operations uses several techniques presented in earlier sections; see App. E for details. We represent RMWAcq assertions analogously to our encoding of Acq assertions (see Sec. 3). We use the value of field `acq` (cf. Fig. 7) to differentiate between holding some RMWAcq assertion

[7] FSL++ also includes a fifth rule for learning that, in certain cases, a CAS operation is *guaranteed* to fail. We have not yet modelled this rule directly (which was not needed for our examples), but believe the extension to be fairly simple if required. Similarly, the general proof rules include optional parameterisation by additional quantifiers, which can be handled, but would complicate the presentation.

from some Acq assertion. Since RMWAcq assertions are duplicable (cf. Fig. 10), we employ `wildcard` permissions for the corresponding `AcqConjunct` predicates.

Our encoding of the proof rule for CAS operations is somewhat involved; we give a high-level description of our approach here, and relegate the details to App. E. We focus on the more-interesting case of a successful CAS operation here. The key challenge is how to select an appropriate assertion T to satisfy the premises of the rule; while we could take T as an annotation, we observed that choosing this overlap to be as large as possible is desirable in practice (since this reduces the resources to be transferred, and which must interact in some cases with the modalities). We write out Viper code to indirectly *compute* this largest-possible T as follows (see Fig. 11(ii) for an illustration).

We introduce yet another heap ("tmp") in which we inhale the invariant for the value read ($\mathcal{Q}[e/\mathcal{V}]$). Now, we attempt to exhale the invariant for the value written ($\mathcal{Q}[e'/\mathcal{V}]$), but adapt the assertions as follows: for each permission in the invariant, we take *the maximum possible* amount from our "tmp" heap; these permissions correspond to T. Any remainder is taken from the current heap (either the real or the "up" heap, depending on τ); these correspond to P. We can express this by rewriting the assertion: for example, if $\tau = \texttt{rlx}$ and $[\![\mathcal{Q}[e'/\mathcal{V}]]\!] \equiv \texttt{acc}(x.f, p)$, we exhale $\mathit{let}\ p' = \mathit{min}(\texttt{perm}(\texttt{tmp}(x).f), p)\ \mathit{in}\ \texttt{acc}(\texttt{tmp}(x).f, p')\ \texttt{\&\&}\ \texttt{acc}(\texttt{up}(x).f, p - p')$. If this modified exhale succeeds, any *remaining* permissions in the "tmp" heap correspond to the assertion A and are moved (in a way similar to our $\texttt{fence}_{\texttt{acq}}$ encoding in Fig. 9) to either the real or "down" heap (depending on τ).

This combination of techniques results in an automatic support for CAS statements and their corresponding proof rule: all of the above steps can be statically generated. This completes the core of our Viper encoding, which now handles the complete set of memory access constructs from Fig. 1.

6 Additional Features

Our encodings support a number of additional features which, for space reasons, we describe only briefly here. Firstly, FSL++ adds rules for *ghost state* to the program logics; ghost locations are fictitious (non-atomic) heap locations added to a program purely for reasoning purposes. For ghost locations, the modalities *collapse*: if an assertion A depends only on ghost locations, then $\triangle A \equiv A \equiv \triangledown A$ holds [12]. Extending our encoding is simple; we track ghost status with an extra boolean function, and tweak our axiomatisation of multiple heaps (Sec. 4) to make the `up` and `down` functions act as identities on ghost locations (details in App. D).

Atomic read and CAS operations are often employed in simple spin loops which wait to read a particular value. Our encoding requires loops to be annotated with invariants in the general case; however, for spin loops (i.e. loops which poll a location via atomic reads or CAS operations until a certain value is read) we can avoid this, based on the following observation: every loop iteration except for the very last (and the first, for atomic reads rather than CAS operations) will make no change to the owned resources. For verification purposes, therefore, it is

sufficient to encode *only* these iterations of spin loops, removing the need for a loop and invariant at the Viper level.

Our encoding also supports fetch-update instructions (such as atomic increment) commands natively. These differ from CAS operations in two ways: they never fail, and the updated value can be a function of the value read. Such commands could also be implemented with CAS operations and loops [12], but this is less efficient for architectures which support atomic fetch-update instructions natively.

The RSL logics do not allow resources to be directly *leaked*; there are no rules for weakening an assertion in order to drop unneeded permissions or assertions such as Rel and Acq assertions during a proof. Nevertheless, the RSL logics do not guarantee absence of memory leaks in general, due to the rules for atomics; there is no guarantee that resources transferred by one thread will be received by another. In order to get rid of unneeded resources, postconditions proved in the RSL logics typically include true as a conjunct, which can be used to hide unwanted assertions. (In such separation logics, true is used to represent arbitrary resources, while emp represents no resources at all). We reflect this practice by adapting the RSL rules to use true instead of emp throughout. In Viper, resources can be leaked by default, at any time. However, using Viper's **perm** construct, we can nonetheless encode additional checks that e.g. ownership of non-atomic locations is never leaked (we do this in our examples).

7 Examples and Evaluation

We have evaluated our encoding by applying it to a number of examples (described below). We are in the process of developing a front-end to generate the appropriate Viper code; for the present time, we have done so by hand. The encoded examples are available in the online appendix [3], and examples of the corresponding input source code (in the language of this paper) are given in App. B.

The encoded versions of the examples make heavy use of Viper macros, to aid readability. Each example consists of three parts: a set of background definitions (e.g. macros to encode each statement), a set of invariant definitions (specific to the example; these reflect the image of the indexing discussed in Sec. 3, along with, e.g., their $\lceil . \rceil^{down}$ variants), and a set of method definitions, corresponding to the translation of the original source code. The three parts are presented in reverse order (so that the source code translation comes first in each file). We also constructed some variations which exercise, e.g., the support for rewriting Acq invariants, since in all other examples our automatic support for tracking conjuncts was sufficient.

As a starting point, we took the examples from the RSL [28] and FSL [13] papers, along with variants of these to demonstrate and test additional features of the logics and our encoding. We also investigated encoding the main example from the latest FSL++ paper [12], which is the Rust reference-counting library [1]. Our encoding does not yet support all of the features of the latest paper; in particular, the proof in [12] requires customising the underlying permission model,

which Viper does not directly support. Following the suggestion of one of the authors [27], we were however able to encode two variants of the original example, in which some of the access modes are strengthened, making the code slightly less efficient but enabling a proof using a simpler permission model. These variants still require *counting permissions* [6], which we were able to express with additional background definitions (reusable for other examples) in the Viper file.

Finally, we tackled seven functions taken from a reader-writer-spinlock in the Facebook Folly library [2]. We were able to verify five of the functions simply, but two employ code idioms which seem to be beyond the scope of the RSL logics, at least without sophisticated extra ghost state; the code depends on the combined effects of multiple atomic accesses (rather than reasoning about these individually). For both functions, we developed alternative implementations which could be handled by the logics and our encoding.

Our initial experiences with our encoding are promising; verification times were reasonable (generally around 10 seconds, and always under a minute), and several errors in our initial specification attempts (and typos in the example code) were uncovered during their development. It is clear that a front-end for our encoding will be necessary to evaluate it on larger codebases; this is largely a matter of engineering (and good error reporting!), and is already underway. The Rust and Facebook code quickly demonstrated a key advantage of working in a general verification framework such as Viper; both examples required support for extra theories (counting permissions as well as modulo and bitwise arithmetic), and we were able to plug such reasoning in simply.

In terms of annotation overhead, the examples require specifications (the ratio between specifications and code was roughly 1:1 for each example tackled so far), but no manual assertions (e.g., between statements) were necessary to make the verification go through. The overhead involved compares favourably to a manual proof effort; an informal proof outline would also require (at least) one intermediate assertion between every statement, while the Coq-formalised proofs of the first RSL examples require roughly ten lines of proof per line of code [28], and for the Rust reference counting example (albeit for an *unmodified* version of the implementation; i.e., a harder proof), the ratio is around 100:1. Such ratios are consistent with other recent Coq-mechanised proofs based on separation logic [29].

8 Conclusions and Future Work

We have presented the first encoding of modern program logics for weak memory models into a deductive program verifier. The encoding enables programs (with suitable annotations) to be verified automatically by existing back-end tools. Moreover, the encoding into a general-purpose verification framework such as Viper provides additional features such as theory reasoning in a straightforward manner (an example being the counting permissions used in our reference counting proofs).

Building front-end verifiers for such advanced logics will enable them to be applied to more, and larger, examples; we are currently working on building a front-end for the translation presented here. In future work, we plan to prove the soundness of the presented encoding; a key ingredient will be the mapping back from a Viper program state to the analogous program state in the RSL logics.

We are interested in tackling other modern program logics such as GPS, which also handles weak memory. We believe that many of the techniques described in this paper can be carried over, and applied to build practical tools that implement these advanced logics. The quick feedback provided by such tools will also enable us to identify weaknesses in the logics themselves, inspiring further developments.

Acknowledgements. We are very grateful to Viktor Vafeiadis and Marko Doko for their explanations of the RSL logics and discussions about our encoding. We also thank Andrei Dan, Lucas Brutschy and Malte Schwerhoff for useful feedback on earlier versions of this manuscript.

References

1. ARC (Atomic Reference Counting) Rust library. Available at `https://doc.rust-lang.org/std/sync/struct.Arc.html`.
2. Facebook Folly reader-writer spinlock implementation. Available at `https://github.com/facebook/folly/blob/master/folly/RWSpinLock.h`.
3. Online appendix. Available at `http://viper.ethz.ch/onlineappendix-rsl-encoding/`.
4. M. Barnett, B.-Y. E. Chang, R. DeLine, B. Jacobs, and K. R. M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *Proceedings of the 4th International Conference on Formal Methods for Components and Objects*, FMCO'05, pages 364–387, Berlin, Heidelberg, 2006. Springer-Verlag.
5. F. Bobot, J.-C. Filliâtre, C. Marché, and A. Paskevich. Why3: Shepherd your herd of provers. In *Boogie 2011: First International Workshop on Intermediate Verification Languages*, pages 53–64, Wrocław, Poland, August 2011.
6. R. Bornat, C. Calcagno, P. O'Hearn, and M. Parkinson. Permission accounting in separation logic. In *Proceedings of POPL '05*, pages 259–270, New York, NY, USA, 2005. ACM.
7. J. Boyland. Checking interference with fractional permissions. In *SAS*, volume 2694 of *LNCS*, pages 55–72. Springer, 2003.
8. A. Chlipala. Mostly-automated verification of low-level programs in computational separation logic. In *Proceedings of the 32Nd ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '11, pages 234–245, New York, NY, USA, 2011. ACM.
9. E. Cohen, M. Dahlweid, M. Hillebrand, D. Leinenbach, M. Moskal, T. Santen, W. Schulte, and S. Tobies. *VCC: A Practical System for Verifying Concurrent C*, pages 23–42. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
10. P. Cuoq, F. Kirchner, N. Kosmatov, V. Prevosto, J. Signoles, and B. Yakobowski. Frama-C. In G. Eleftherakis, M. Hinchey, and M. Holcombe, editors, *Software Engineering and Formal Methods: 10th International Conference, SEFM 2012, Thessaloniki, Greece, October 1-5, 2012. Proceedings*, pages 233–247, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

11. T. Dinsdale-Young, P. D. R. Pinto, K. J. Andersen, and L. Birkedal. Caper: Automatic verification for fine-grained concurrency. *European Symposium on Programming (ESOP) 2017*. To appear. Draft available from `http://www.cs.au.dk/~{}birke/papers/`.
12. M. Doko and V. Vafeiadis. Tackling real-life relaxed concurrency with FSL++. *European Symposium on Programming (ESOP) 2017*. To appear. Draft available from `http://plv.mpi-sws.org/fsl/`.
13. M. Doko and V. Vafeiadis. A program logic for C11 memory fences. In *VMCAI*, volume 9583 of *Lecture Notes in Computer Science*, pages 413–430. Springer, 2016.
14. S. Heule, K. R. M. Leino, P. Müller, and A. J. Summers. Abstract read permissions: Fractional permissions without the fractions. In *Verification, Model Checking, and Abstract Interpretation (VMCAI)*, volume 7737 of *Lecture Notes in Computer Science*, pages 315–334. Springer-Verlag, 2013.
15. B. Jacobs. Verifying TSO programs. CW Reports CW660, Department of Computer Science, KU Leuven, May 2014.
16. K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. In *Proceedings of LPAR'10*, pages 348–370, Berlin, Heidelberg, 2010. Springer-Verlag.
17. K. R. M. Leino and P. Müller. A basis for verifying multi-threaded programs. In G. Castagna, editor, *ESOP*, volume 5502 of *LNCS*, pages 378–393. Springer-Verlag, 2009.
18. P. Müller, M. Schwerhoff, and A. J. Summers. Automatic verification of iterated separating conjunctions using symbolic execution. In S. Chaudhuri and A. Farzan, editors, *Computer Aided Verification (CAV)*, volume 9779 of *LNCS*, pages 405–425. Springer-Verlag, 2016.
19. P. Müller, M. Schwerhoff, and A. J. Summers. Viper: A verification infrastructure for permission-based reasoning. In B. Jobstmann and K. R. M. Leino, editors, *VMCAI*, volume 9583 of *LNCS*, pages 41–62. Springer-Verlag, 2016.
20. P. W. O'Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *Proceedings of CSL'01*, pages 1–19, London, UK, UK, 2001. Springer-Verlag.
21. M. J. Parkinson and A. J. Summers. The relationship between separation logic and implicit dynamic frames. *Logical Methods in Computer Science*, 8(3:01):1–54, 2012.
22. J. C. Reynolds. Definitional interpreters for higher-order programming languages. In *ACM Annual Conference—Volume 2*, ACM '72, pages 717–740. ACM, 1972.
23. J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*. IEEE Computer Society Press, 2002.
24. I. Sergey, A. Nanevski, and A. Banerjee. Mechanized verification of fine-grained concurrent programs. In *Proceedings of PLDI '15*, pages 77–87, New York, NY, USA, 2015. ACM.
25. J. Smans, B. Jacobs, and F. Piessens. Implicit dynamic frames. *ACM Trans. Program. Lang. Syst.*, 34(1):2:1–2:58, 2012.
26. A. Turon, V. Vafeiadis, and D. Dreyer. GPS: navigating weak memory with ghosts, protocols, and separation. In *OOPSLA*, pages 691–707. ACM, 2014.
27. V. Vafeiadis. Personal communication, December 2015.
28. V. Vafeiadis and C. Narayan. Relaxed separation logic: a program logic for C11 concurrency. In *OOPSLA*, pages 867–884. ACM, 2013.
29. F. Xu, M. Fu, X. Feng, X. Zhang, H. Zhang, and Z. Li. A practical verification framework for preemptive OS kernels. In S. Chaudhuri and A. Farzan, editors, *Computer Aided Verification: 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part II*, pages 59–79, Cham, 2016. Springer International Publishing.

A Full Source Encoding

The encoding of the general source language of assertions is given below (we assume the encoding of pure expressions, which can typically be the identity mapping, assuming all operators supported such as addition, equality etc. are all supported natively by Viper).

$$
\begin{aligned}
\llbracket \mathsf{emp} \rrbracket &\rightsquigarrow \mathsf{true} \\
\llbracket l \overset{k}{\mapsto} e \rrbracket &\rightsquigarrow \texttt{acc(l.val, k) \&\& acc(l.init, k) \&\& l.val==}\llbracket e \rrbracket \texttt{ \&\& l.init} \\
\llbracket A_1 * A_2 \rrbracket &\rightsquigarrow \llbracket A_1 \rrbracket \texttt{\&\&} \llbracket A_2 \rrbracket \\
\llbracket b \Rightarrow A \rrbracket &\rightsquigarrow \llbracket b \rrbracket \Rightarrow \llbracket A \rrbracket \\
\llbracket (b \ ?\ A_1 : A_2) \rrbracket &\rightsquigarrow (\llbracket b \rrbracket \ ?\ \llbracket A_1 \rrbracket : \llbracket A_2 \rrbracket) \\
\llbracket \mathsf{Uninit}(l) \rrbracket &\rightsquigarrow \texttt{acc(l.val) \&\& acc(l.init) \&\& !l.init} \\
\llbracket \mathsf{Acq}(l, \mathcal{Q}) \rrbracket &\rightsquigarrow \texttt{acc(l.acq, wildcard) \&\& l.acq==true \&\}} \\
&\quad (\textit{foreach } i \textit{ in } \langle\!\langle \mathcal{Q} \rangle\!\rangle \texttt{: AcqConjunct(l, } i\texttt{) \&\&} \\
&\qquad \texttt{valsRead(l, } i\texttt{)==Set[Int]()} \textit{ end}) \\
\llbracket \mathsf{Rel}(l, \mathcal{Q}) \rrbracket &\rightsquigarrow \texttt{acc(l.rel, wildcard) \&\& l.rel==}\langle \mathcal{Q} \rangle \\
\llbracket \mathsf{Init}(l) \rrbracket &\rightsquigarrow \texttt{acc(l.init, wildcard)} \\
\llbracket \triangle A \rrbracket &\rightsquigarrow \lceil \llbracket A \rrbracket \rceil^{up} \\
\llbracket \nabla A \rrbracket &\rightsquigarrow \lceil \llbracket A \rrbracket \rceil^{down} \\
\llbracket \mathsf{RMWAcq}(l, \mathcal{Q}) \rrbracket &\rightsquigarrow \texttt{acc(l.acq, wildcard) \&\& l.acq==false \&\}} \\
&\quad (\textit{foreach } i \textit{ in } \langle\!\langle \mathcal{Q} \rangle\!\rangle \texttt{: acc(AcqConjunct(l, } i\texttt{), wildcard)} \textit{ end})
\end{aligned}
$$

B Example Details

To give an impression of the input required for our encoding, we provide source code corresponding to some of our encoded examples from the Online Appendix [3], given in Figures 12 to 16.

C Rewriting Invariants

It is unusual (at least, in the examples we have investigated so far) for very many different invariants for atomic locations to be needed; it is even less common for there to be a need for many different invariants for the *same* atomic location. Indeed, for Rel and RMWAcq assertions, since the assertions are duplicable, one may as well always use the same invariant. For Acq assertions the situation is more interesting; it may be desirable to split the invariant (as used e.g. in Fig. 15) across several Acq assertions, and programmer-annotated assertions may not always syntactically match up precisely (since there might be more readable ways of expressing an equivalent assertion). Since our indexing of Acq invariants matches their conjuncts syntactically, additional work is required if this syntactic match would be overly restrictive. For example, in the example shown in Fig. 15, the initial Acq invariant is expressed more succinctly in a way which does not provide the immediate conjuncts needed by the (left and

$\mathcal{Q} \equiv (\mathcal{V} = 0 \ ? \ \mathsf{true} : (\mathcal{V} = 1 \ ? \ J : \mathsf{false}))$

$Lock(x) \equiv \mathsf{Init}(x) * \mathsf{RMWAcq}(x, \mathcal{Q}) * \mathsf{Rel}(x, \mathcal{Q})$

```
new_lock() returns (x)
  requires J
  ensures Lock(x)
{
  x := alloc_acq(Q);
  [x]_rel := 1;
}

lock(x)
  requires Lock(x)
  ensures J * Lock(x)
{
  while(CAS_rel_acq(x, 1, 0) != 1);
}

unlock(x)
  requires J * Lock(x)
  ensures Lock(x)
{
  [x]_rel := 1;
}
```

Fig. 12. RSL Figure 7 example, written in our source language. Annotations in blue.

$\mathcal{Q} \equiv (\mathcal{V} = 0 \ ? \ \mathsf{true} : a \overset{1}{\mapsto} 7)$

$\{\mathsf{emp}\}$

$a := \mathtt{alloc}_{\mathsf{na}}();$

$c := \mathtt{alloc}_{\mathsf{acq}}(\mathcal{Q});$

$[c]_{\mathsf{rlx}} := 1;$

$\{\mathsf{Uninit}(a) * \mathsf{Rel}(c, \mathcal{Q})\}$	$\{\mathsf{Acq}(c, \mathcal{Q}) * \mathsf{Init}(c)\}$
$[a]_{\mathsf{na}} := 7$	$\mathtt{while}([c]_{\mathsf{acq}} == 0);$
$[c]_{\mathsf{rel}} := 1$	$x := [a]_{\mathsf{na}}$
$\{\mathsf{Rel}(c, \mathcal{Q})\}$	$[a]_{\mathsf{na}} := x + 1$
	$\{a \overset{1}{\mapsto} 8 * \mathsf{true}\}$

$\{a \overset{1}{\mapsto} 8 * \mathsf{true}\}$

Fig. 13. RSL Figure 8 example, written in our source language. Annotations in blue.

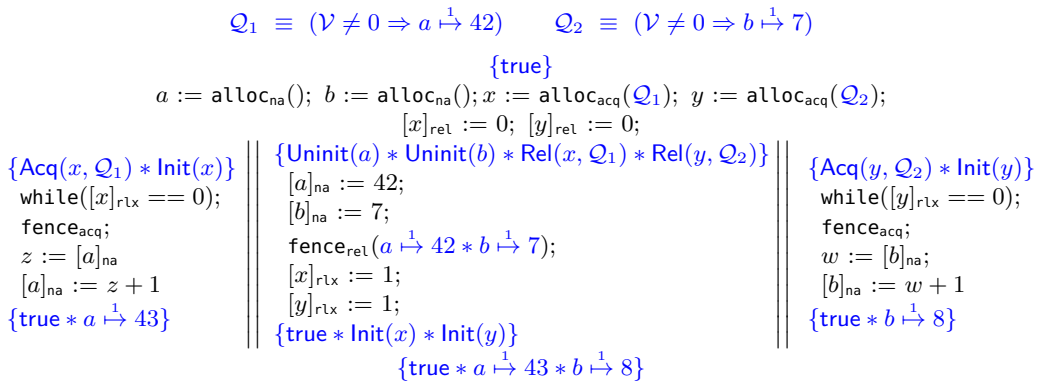

Fig. 14. FSL Figure 2 example, written in our source language. Annotations in blue.

right) forked threads. In such cases, we support an additional `rewrite` statement `rewrite` $\mathsf{Acq}(l, \mathcal{Q})$ `as` $\mathsf{Acq}(l, \mathcal{Q}')$ in the source program to explain the intended rewriting. To *check* that such a statement is justified, we need to check the entailment between the original and new invariants, for all values of $\mathcal{V}$. Furthermore, this entailment check cannot be made directly in the *current* state, since that might already contain permissions and value information which could unsoundly weaken the check made, or even contradict the invariants involved, resulting in an infeasible program state from there onwards.

To avoid these issues, we perform the following steps (shown in Fig. 17). For simplicity, we do not handle the case of rewriting invariants for which values have already been read (we check that this is not the case, here, but an extension is possible). Firstly, we create a non-deterministic **`if`**-branch. Inside the branch we *remove* all permissions from the current state. We then havoc an integer variable, representing the arbitrary value of $\mathcal{V}$ for which the subsequent check should hold. We **`inhale`** the original invariant (using our indexing as usual), and **`exhale`** the invariant to which it is to be rewritten. If this succeeds, we know that the rewriting is justified; the one invariant entails the other, for all values of $\mathcal{V}$. We then kill this branch, by adding an **`assume false`** to it; subsequently, only the other branch (in which no changes were made) will be considered for verification.

Lastly, we perform the rewriting itself by discarding all of the original `AcqConjunct` instances, and replacing them with the new ones. Verification can then proceed as usual.

D Ghost Locations

We extend our encoding to handle ghost locations in a simple manner. Firstly, we add a boolean function `is_ghost` on references, to identify whether or not a location is ghost. We added macros `realRef(r)` to add the appropriate assumptions for a real program reference, and `ghostRef(r)` for ghost locations; this can be seen

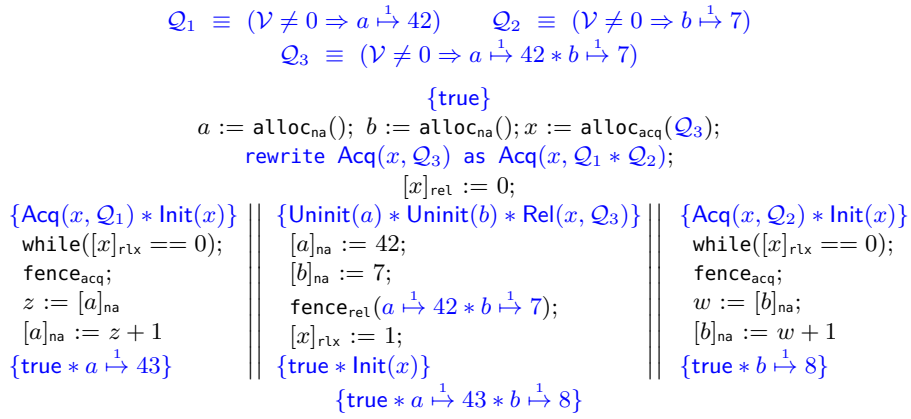

Fig. 15. Variant of FSL Figure 2 example (using invariant rewriting), written in our source language. Annotations in blue.

as the translation of "type information", since we assume that ghost locations are labelled as such in the original program.

For ghost locations we tweak our multiple heaps encoding to change the assumptions about the `up` and `down` mappings to instead require them to act as the identity (correspondingly, the result of `heap` is no longer constrained to be different after applying these mappings to a ghost reference). This immediately gives us that, for assertions depending only on ghost locations in the heap, $\triangle A$, A and $\triangledown A$ will be handled equivalently; since (up to syntactic applications of these identify mappings) they will be encoded as identical assertions.

Finally, we add an assumption of `realRef(r)` to our existing statements for allocating references, and add a new ghost allocation statement, for which the analogous `ghostRef(r)` assumption is added. The most-relevant details are summarised in Fig. 18.

E Compare and Swap Details

The details of our encoding of the FSL++ compare and swap rules (cf. Fig. 10) are shown in Fig. 19. We represent RMWAcq assertions similarly to Acq assertions (cf. Fig. 7), but and a `false` value of the `acq` field to differentiate holding one from the other. Recall that we must choose at allocation whether atomic reads or compare and swaps will be used to gain ownership via the atomic location; this choice is then reflected in the field value. The encoding of allocation is then straightforward.

The handling of a CAS statement itself involves initially checking that we indeed hold some Init, RMWAcq and Rel() assertions for the location, according to the precondition of the rule, and then using an if-condition over the fresh read value `x` to narrow us down to the case of a successful CAS. The subsequent

$\mathcal{Q} \;\equiv\; \mathcal{V} \geq 0 * g \overset{1-\mathcal{V}*rd}{\mapsto} _ * (\mathcal{V} \geq 1 \Rightarrow d \overset{1-\mathcal{V}*rd}{\mapsto} _)$

$ARC(d,c,g,v) \;\equiv\; d \overset{rd}{\mapsto} v * g \overset{rd}{\mapsto} _ * \mathsf{RMWAcq}(c,\mathcal{Q}) * \mathsf{Rel}(c,\mathcal{Q}) * \mathsf{Init}(c)$

```
new(v) returns (d,c,g)
  requires true
  ensures ARC(d,c,g,v)
{
  d := alloc_na();
  g := alloc_ghost();
  c := alloc_RMW(Q);
  [d]_na := v;
  [c]_rel := 1;
}

read(d,c,g) returns (v)
  requires ARC(d,c,g,_)
  ensures ARC(d,c,g,v)
{
  v := [d]_na;
}

drop(d,c,g)
  requires ARC(d,c,g,_)
  ensures true
{
  t := fetch_and_add_rel(c,-1);
  if (t==1){
    fence_acq;
    free(d);
  }
}

clone(d,c,g)
  requires ARC(d,c,g,v)
  ensures ARC(d,c,g,v) * ARC(d,c,g,v)
{
  fetch_and_add_acq(c,1);
}
```

Fig. 16. Rust reference counting variant 2 (in online appendix). Compared to the original code (see [12]) we modified the write in `new` to use a release rather than relaxed mode, and the update in `clone` to use acquire rather than relaxed. As discussed in the paper body, the original version of the example is proved in [12] using features which are not yet supported by our encoding. We do, however, exploit the CAS rules and rules for fences here. We write *rd* for a read permission, in the sense of counting permissions [6]. The ghost location g must be identifiable as such for the encoding, for example by considering this a type annotation, or using a distinguished class of variables for ghost locations. We model the `free` statement by exhaling the corresponding permissions.

```
⟦rewrite Acq(l, Q) as Acq(l, Q')⟧ ⇝
  assert SomeAcq(l)
  var tmpBool : Bool
  tmpBool := havoc()

  if(tmpBool) { // check rewriting is justified

// remove all permissions from current state
    exhale forall r: Ref :: r != null ==> acc(r.init, perm(r.init))
    exhale forall r: Ref :: r != null ==> acc(r.val, perm(r.val))
    exhale forall r: Ref :: r != null ==> acc(r.rel, perm(r.rel))
    exhale forall r: Ref :: r != null ==> acc(r.acq, perm(r.acq))
// analogously for other fields, predicates in source program

    var v :Int
    v := havoc() // perform check for arbitrary v

// inhale original invariant
    foreach i in indices do
    if(i in ⟪Q⟫) {
      inhale inv(i)[v/V]
    }
    end

// exhale new invariant
    foreach i in indices do
    if(i in ⟪Q'⟫) {
      exhale inv(i)[v/V]
    }
    end

    assume false // kill this branch - we've checked rewriting is OK
  }

// update the conjuncts held
  exhale (foreach i in ⟪Q⟫:
  AcqConjunct(l, i) && valsRead(l, i) == Set[Int]() end)
  inhale (foreach i in ⟪Q'⟫:
  AcqConjunct(l, i) && valsRead(l, i) == Set[Int]() end)
```

Fig. 17. Viper encoding of a source-level `Rewrite` statement.

Viper code reflects the three steps described in Sec. 5 and Fig. 11. Firstly, we perform the inhale of newly-gained resources (corresponding to $\mathcal{Q}[e/\mathcal{V}]$) into the `tmp` heap.

Secondly, we attempt to exhale the assertion $\mathcal{Q}[e'/\mathcal{V}]$, modified so that the permissions are taken preferentially from the `tmp` heap, and failing this, from the real heap or `up` heaps, depending on whether or not the write synchronises. This modification of the assertion (which splits the permission amounts across the two heaps, as described in Sec. 5) is denoted by the $\lceil . \rceil^{tmp/real}$ and $\lceil . \rceil^{tmp/up}$ mappings; if the *values* of heap locations are also mentioned in the parameter assertions, then these heap dereferences must also be rewritten to a dereference in the corresponding heap (e.g. `x.val == 4` might become `tmp`(x)`.val == 4`). In case permission to the corresponding location is taken partly from both heaps,

```
define realRef(x) !is_ghost(x) && heap(x) == 0
define ghostRef(x) is_ghost(x) && heap(x) == 0

domain parallelHeaps {
  function up(x: Ref) : Ref
  function down(x: Ref) : Ref
  function up_inv(x: Ref) : Ref
  function down_inv(x: Ref) : Ref

  function temp(x: Ref) : Ref
  function temp_inv(x: Ref) : Ref

  function heap(x: Ref) : Int
  function is_ghost(x:Ref) : Bool

  axiom { forall r:Ref :: up_inv(up(r)) == r &&
    (is_ghost(r) ? up(r) == r : heap(r)==0 ==> heap(up(r)) == 1) }
  axiom { forall r:Ref :: {up_inv(r)} up(up_inv(r)) == r &&
    (is_ghost(r) ? up_inv(r) == r : heap(r)==1 ==> heap(up_inv(r)) == 0) }
  axiom { forall r:Ref :: {down(r)} down_inv(down(r)) == r  &&
    (is_ghost(r) ? down(r) == r : heap(r)==0 ==> heap(down(r)) == -1) }
  axiom { forall r:Ref :: {down_inv(r)} down(down_inv(r)) == r &&
    (is_ghost(r) ? down_inv(r) == r : heap(r)==-1 ==> heap(down_inv(r)) == 0) }

  axiom { forall r:Ref :: {temp(r)} temp_inv(temp(r)) == r  &&
    (is_ghost(r) ? temp(r) == r : heap(r)==0 ==> heap(temp(r)) == -3) }
  axiom { forall r:Ref :: {temp_inv(r)} temp(temp_inv(r)) == r &&
    (is_ghost(r) ? temp_inv(r) == r : heap(r)==-3 ==> heap(temp_inv(r)) == 0) }
}

⟦l := alloc_ghost()⟧ ⇝
  x := new(); assume ghostRef(x); // ghost location
  inhale Uninit(x) // ghost locations are non-atomics
```

Fig. 18. Extension of our Viper encoding to handle ghost locations.

the extra assumption that the two values are the same can be explicitly added by these mappings.

Finally (assuming the exhale has succeeded, otherwise a verification failure will have been encountered), all remaining permissions in the `tmp` heap are transferred to either the real or `down` heap, depending on whether the read synchronises.

```
define SomeRMWAcq(l) acc(l.acq, wildcard) && l.acq == false

⟦RMWAcq(l, Q)⟧ ⇝ SomeRMWAcq(l) &&
  (foreach i in ⟪Q⟫: acc(AcqConjunct(l, i),wildcard) end)

⟦l := alloc_RMW(Q)⟧ ⇝
  x := new(); assume realRef(x); // not a ghost location
  inhale ⟦Rel(l, Q)⟧ && ⟦RMWAcq(l, Q)⟧

⟦x := CAS_τ(l, e, e′)⟧ ⇝
  assert Init(l) && SomeRMWAcq(l) && SomeRel(l)
  x := havoc()
// inhale into tmp heap
  if(x == ⟦e⟧) {
    foreach i in indices do
      if (perm(AcqConjunct(l, i)) > 0) {
        inhale ⌈inv(i)⌉^tmp[x/V]
      }
    end
// exhale from tmp &&real/up heaps (depending on τ)
    foreach i in indices do
      if (i == l.rel) { // write synchronises
        if (τ ∈ {rel, rel_acq}) {
          exhale ⌈inv(i)⌉^tmp/real[⟦e′⟧/V]
        } else {
          exhale ⌈inv(i)⌉^tmp/up[⟦e′⟧/V]
        }
      }
    end
// ... move tmp heap to real/down heap (depending on τ)
    var rs : Set[Ref]; rs := havoc() // unknown set of Refs
    assume forall r: Ref :: r in rs <==> perm(tmp(r).val) > none
    if(τ ∈ {acq, rel_acq}) {
      inhale forall r: Ref :: r in rs ==> acc(r.val, perm(tmp(r).val))
      assume forall r: Ref :: r in rs ==> r.val == tmp(r).val
    } else {
      inhale forall r: Ref :: r in rs ==> acc(down(r).val, perm(tmp(r).val))
      assume forall r: Ref :: r in rs ==> down(r).val == tmp(r).val
    }
    exhale forall r: Ref :: r in rs ==> acc(tmp(r).val, perm(tmp(r).val))
// analogously for each other field, predicate (in place of val)
  }
```

Fig. 19. Viper encoding of the RSL rules for compare and swap operations.